



# DISCOM2 Networking at SC99

## A Step on the Path to 100 Gigabit Per Second

The networking demonstration within the ASCI booth is designed to highlight and demonstrate the latest available networking technologies while serving the booth's communication needs. The Distance Computing and Distributed Computing Program (DisCom<sup>2</sup>) provides the network design for the ASCI booth. DisCom<sup>2</sup> is intended to deliver key computing and communications technologies that enable the ASCI vision. DisCom<sup>2</sup> will provide the technology solution to efficiently integrate distributed resources with high-end computing resources at a distance. DisCom<sup>2</sup> will use the SC99 forum to validate the performance and capabilities of new communication technologies while increasing the understanding of the high performance networking technologies available to the ASCI communities.

The networking demonstration within the ASCI booth consists of many large switching components that are grouped by core technology into three networks. Each network individually constitutes a high performance network. The three networks are designated by the technology being used at their core, Gigabit Ethernet, ATM, and Packet Over SONET (POS). Each network's core elements are capable of peak performances ranging from 20 gigabits per second to terabits per second and contain external connections equal to or greater than 1 Gigabit per second. The total external bandwidth available from the booth is greater than 6 gigabits per second. The three technology networks are interconnected and interoperating in the booth to emphasize the importance of the interoperation of these high performance, networking technologies.

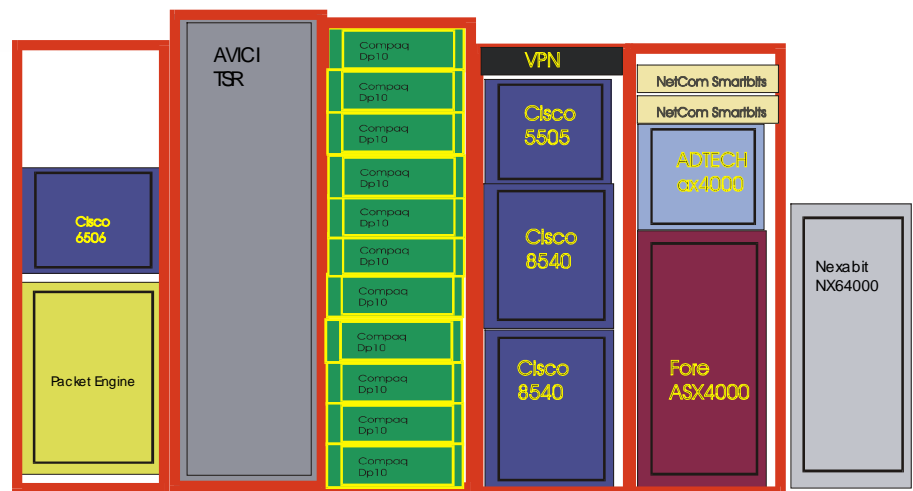


Figure 1. Core Networking Equipment at SC99

### Network Equipment

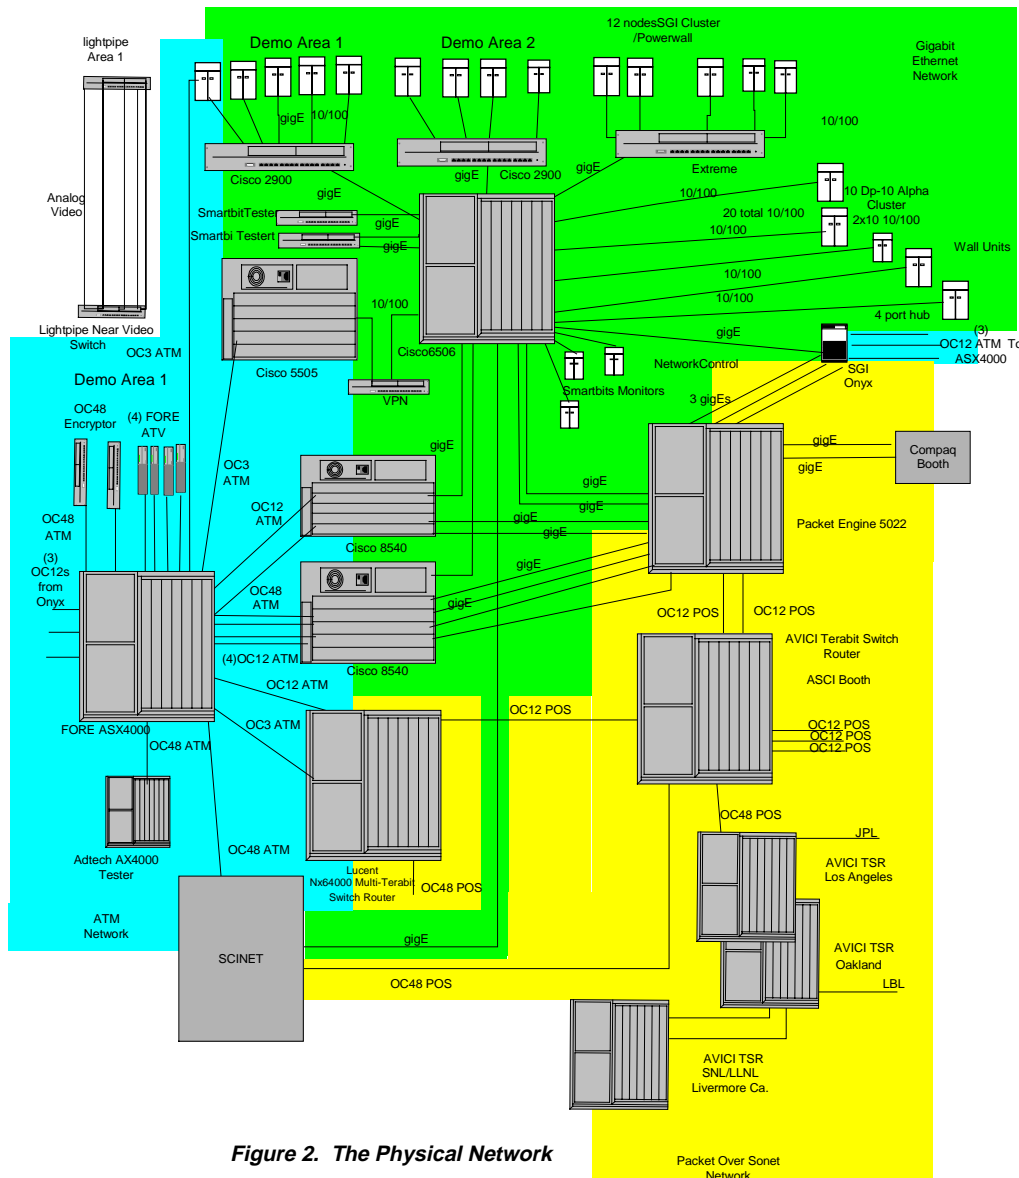
Starting on the left side of the exhibit is a Cisco Systems 6506 Ethernet Switch, which is the core element of the booth's Gigabit Ethernet network. The Cisco 6506 is capable of 30 gigabits per second but only supports Ethernet interfaces at this time. Below the 6506 is the Packet Engines 5200, the interconnection point between the Gigabit Ethernet network and the POS network. The Packet Engines 5200 also provides the connection between the Compaq compute cluster, the SGI Video Powerwall, and the computation cluster located in the COMPAQ booth, via the experimental network (XNET). The Packet Engine 5200 is capable of 52 gigabits per second and currently supports Ethernet and POS interfaces. The AVICI Terabit Switch Router (TSR), located in the next cabinet is the core of the POS network and an external entry point for SC99's XNET. The TSR is capable of scaling to 1000s of gigabit interfaces and currently supports OC48 and OC12 POS interfaces. In the next cabinet is

a Compaq cluster consisting of ten DP10 workstations interconnected to the booth network by (20) 100 megabit Ethernet interfaces. This booth cluster is connected to similar clusters at Sandia National Laboratories in Livermore, California and the computation cluster in the Compaq booth. These interconnected clusters allow us to evaluate the performance of computations executed on machines that are local and remote. The next cabinet contains three network switches, two Cisco's 8540s at the bottom of the rack and a Cisco 5505. The two Cisco 8540s will pass data from the large computing platforms at Lawrence Livermore National Laboratory and Los Alamos Laboratory to the DISCOM Wide Area Network. The 8540s support Gigabit Ethernet and ATM Interfaces and are rated at 20 Gigabit per second switches. The Cisco 5505 supports a Virtual Private Network (VPN) within the booth. And provides a path between 10/100 Ethernet and the ATM network. The Cisco 5505 supports 10/100baseT Ethernet, OC3, and OC12 ATM interfaces and is rated at 3 gigabits per second. At the very top of the

rack is a Network Alchemy CryptoCluster 2500, which implements the IPSEC Triple DES encryption and firewall functions that are key elements in a VPN. This VPN provides the necessary security features for a remote visualization demonstration in the booth. The next cabinet contains a Fore

of the ASCII booth's architecture. The cabinet on the far right of the networking area contains a Lucent NX64000 Multi-Terabit Switch/Router. This router currently supports OC192 POS, OC48c POS, OC12c POS, OC3c POS, OC12c ATM, OC3c ATM, and DS3 frame interfaces with line-speed

performance network, computation resources capable of tens of teraOPs, and storage capacity growing to petabytes to the hundreds of scientific and engineering users spread throughout the United States and you have the ASCII network for the year 2000.



**Figure 2. The Physical Network**

System ASX4000 ATM switch which is the core element in the booth's ATM network. The ASX4000 support OC3, OC12, and OC48 ATM interfaces and is rated at 40 gigabits per second. Above the ASX4000 is an Adtech AX4000 Analyzer/Generator that supports POS, ATM and Gigabit Ethernet testing for line interface at OC48, GigE, Oc12, and Oc3 rates. Above the AX4000 are two Netcom Smartbits testers that support GigE, ATM, POS, and 10/100 Ethernet testing. The Smartbits and the AX4000 will be used to test the performance

QOS features. The NX64000 provides the booth with POS to ATM transport. The NX64000 is rated at 6.4 terabits per second.

The ASCII booth network at SC99 is designed to provide multiple gigabit per second networking performance to the booth demonstrations with the largest consumer of bandwidth being the distributed visualization demonstrations running on the SGI Powerwall. Add the booth's high

For more information contact:

Thomas Pratt  
Sandia National Laboratories  
Ms0806  
PO box 5800  
Albuquerque NM 87123  
tjpratt@sandia.gov

